

+

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office, U.S. DEPARTMENT OF COMMERCE

PTO/SB/17 (12/99)

PTO

09/685657

10/10/00

FEE TRANSMITTAL

for FY 2000

Patent fees are subject to annual revision.
Small Entity payments must be supported by a small entity statement,
otherwise large entity fees must be paid. See Forms PTO/SB/09-12.
See 37 C.F.R. §§ 1.27 and 1.28.

TOTAL AMOUNT OF PAYMENT (\$) 828.00

Complete if Known

Application Number
Filing Date October 10, 2000
First Named Inventor Vipul Shah, et al.
Examiner Name
Group / Art Unit
Attorney Docket No. 26530.23(IDR-464/5)

METHOD OF PAYMENT (check one)

1. ☒ The Commissioner is hereby authorized to charge indicated fees and credit any overpayments to:

Deposit Account Number 08-1394

Deposit Account Name Haynes and Boone LLP

☐ Charge Any Additional Fee Required Under 37 CFR §§ 1.16 and 1.17

2. ☒ Payment Enclosed:

☒ Check ☐ Money Order ☐ Other

FEE CALCULATION

1. BASIC FILING FEE

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description	Fee Paid
101 690	201 345	Utility filing fee	710.00
106 310	206 155	Design filing fee	
107 480	207 240	Plant filing fee	
108 690	208 345	Reissue filing fee	
114 150	214 75	Provisional filing fee	

SUBTOTAL (1) (\$) 710.00

2. EXTRA CLAIM FEES

Total Claims 18 - 20** = 0 X Fee from below = 0.00
Independent Claims 4 - 3** = 1 X 78 = 78.00
Multiple Dependent
Fee Paid

**or number previously paid, if greater; For Reissues, see below

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description	Fee Paid
103 18	203 9	Claims in excess of 20	
102 78	202 39	Independent claims in excess of 3	
104 260	204 130	Multiple dependent claim, if not paid	
109 78	209 39	** Reissue independent claims over original patent	
110 18	210 9	** Reissue claims in excess of 20 and over original patent	

SUBTOTAL (2) (\$) 78.00

FEE CALCULATION (continued)

3. ADDITIONAL FEES

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description	Fee Paid
105 130	205 65	Surcharge - late filing fee or oath	
127 50	227 25	Surcharge - late provisional filing fee or cover sheet	
139 130	139 130	Non-English specification	
147 2,520	147 2,520	For filing a request for reexamination	
112 920*	112 920*	Requesting publication of SIR prior to Examiner action	
113 1,840*	113 1,840*	Requesting publication of SIR after Examiner action	
115 110	215 55	Extension for reply within first month	
116 380	216 190	Extension for reply within second month	
117 870	217 435	Extension for reply within third month	
118 1,360	218 680	Extension for reply within fourth month	
128 1,850	228 925	Extension for reply within fifth month	
119 300	219 150	Notice of Appeal	
120 300	220 150	Filing a brief in support of an appeal	
121 260	221 130	Request for oral hearing	
138 1,510	138 1,510	Petition to institute a public use proceeding	
140 110	240 55	Petition to revive - unavoidable	
141 1,210	241 605	Petition to revive - unintentional	
142 1,210	242 605	Utility issue fee (or reissue)	
143 430	243 215	Design issue fee	
144 580	244 290	Plant issue fee	
122 130	122 130	Petitions to the Commissioner	
123 50	123 50	Petitions related to provisional applications	
126 240	126 240	Submission of Information Disclosure Stmt	
581 40	581 40	Recording each patent assignment per property (times number of properties)	40.00
146 690	246 345	Filing a submission after final rejection (37 CFR § 1.129(a))	
149 690	249 345	For each additional invention to be examined (37 CFR § 1.129(b))	

Other fee (specify)

Other fee (specify)

* Reduced by Basic Filing Fee Paid

SUBTOTAL (3) (\$) 40.00

SUBMITTED BY

Name (Print/Type)	David L. McCombs	Registration No. (Attorney/Agent)	32,271	Telephone	(214) 651-5533
Signature		Date	10-10-00		

WARNING:

Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231.

d789702.1

ATTORNEY DOCKET NO.26530.23(IDR-464/5)
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE



Applicant:	Vipul Shah, et al.	§		
Serial No.:	N/A	§	Group Art Unit:	Unknown
Filed:	Herewith	§	Examiner:	Unknown
For:	SYSTEM AND METHOD FOR BALANCING TCP/IP WORKLOAD OF MULTI-PROCESSOR SYSTEM BASED ON HASH BUCKETS	§ § § § § §		

Commissioner for Patents
Box Patent Application
Washington, D.C. 20231

EXPRESS MAIL CERTIFICATE

Express Mail Number: EL418585737US

Date of Deposit: October 10, 2000

I hereby certify that the following attached papers and fee:

1. Utility Patent Application Transmittal and Fee Transmittal with duplicate copy attached;
2. Utility Patent Application consisting of: 16 pages of Specification;
3. 2 Informal Drawing sheets;
4. a signed Declaration;
5. a Check in the amount of \$788.00;
6. an executed Assignment and Assignment Recordation Cover Sheet;
7. a Check in the amount of \$40.00 for Assignment Recordation; and
8. a Return Postcard.

are being deposited with United States Postal Service "Express Mail Post Office to addressee" to the
Assistant Commissioner for Patents, Washington, D. C. 20231.


Debbie Ludwig

10-10-2000
Date

EXPRESS MAIL NO.: EL41858573705 DATE OF DEPOSIT: 10 - 10 2000

This paper and fee are being deposited with the U.S. Postal Service Express Mail Post Office to Addressee service under 37 CFR §1.10 on the date indicated above and is addressed to the Commissioner for Patents, Washington, D.C. 20231

Debbie Ludwig
Name of person mailing paper and fee

Debbie Ludwig
Signature of person mailing paper and fee

SYSTEM AND METHOD FOR BALANCING TCP/IP/ WORKLOAD OF MULTI-PROCESSOR SYSTEM BASED ON HASH BUCKETS

Background of the Invention

The present invention relates generally to computer software, and more particularly, to a system and method for utilizing hash buckets for managing workload in a multiple processor computer system.

5 In today's computer network environment, large volumes of data are customarily shared by multiple computers. In order to improve upon the efficiency of the computer network, multiple computer processors share the workload as well. Especially in a multiple processor computer system where a network protocol such as TCP/IP is used to link all participating processors, it is
10 desirable to have the TCP/IP workload evenly distributed across these processors.

The current art finds some solutions by dealing with the workload distribution problem from three different levels. That is, the solutions are found in methods such as packet level workload distribution, layer level workload
15 distribution, and connection level workload distribution.

According to the packet level distribution method, the loads of the processors at a packet level are monitored and IP packets are distributed across processors accordingly. This method seems to be fine-grained because the granularity of distribution is a single packet. However, this method has been
5 proven to be an audacious proposal and does not work well for connection oriented protocols like TCP (although this technique has been proved to be effective while working with layers like UDP). Anomalies of packet level workload distribution such as out-of-order packet handling can play havoc with network performance and render the packet level workload distribution
10 unreliable.

The layer level distribution method distributes the workload of various layers like IP, TCP, UDP onto various processors. This introduces additional queuing delays and context switches per packet level. Thus, it is also found to be limited in exploiting available parallel processing mechanism. Moreover, this
15 mechanism can not scale to large numbers of processors since there are only a small number of layers.

For the connection level distribution method distributing TCP workload, statically mapping connections to processors can lead to load imbalance, and other negative performance impacts for a specific processor-bound application.
20 The traditional dynamic mapping mechanism is not scalable, since the number of connections can be in the tens of thousands and re-balancing such large numbers of connections cannot be done in a performance-sensitive path like the network protocol path.

The main issue is to ensure that the workload distribution is equitable so
25 that no processor is overloaded or underloaded thereby causing performance bottlenecks or resource waste. On the other side, because the network protocol path is performance sensitive, it is important that the workload balancing

mechanism itself is not process intensive or resource intensive which adds itself as a "heavy weight" load to the processors.

Another technical bottleneck in a multiple processor system is a single timer thread structure, in which only one timer thread is designated for all processors. In order to optimize the advantage of having multiple processors in the system, multiple timer threads are desired to be used for parallel processing.

What is needed is an efficient method and system for equitably distributing process workload in a multiple processor network environment and an appropriate mechanism for using multiple timer threads.

10

Summary of the Invention

A method and system is provided for balancing workload for a plurality of computer processors in a multiple processor computer system. In one embodiment, there are a plurality of packets from different connections needed to be processed. First, a plurality of hash buckets are predetermined by a hash function, and a connection is uniquely mapped to a hash bucket. A load balancer assigns the hash buckets through multiple queues to different processors so that the workload of the processors can be balanced. A queue model may also be applied to assure each processor gives a fair attention to packets from different connections in the hash bucket.

In another embodiment of the present invention, and in connection with processing connections with a timer thread, a method and system is further provided for processing a plurality of connections with a plurality of timer threads. A plurality of hash buckets are provided, and each connection is mapped to one of the hash buckets. Further, each hash bucket is assigned to a processor based on its workload so that the processor processes the connections relating to the assigned hash bucket, wherein the plurality of timer threads thus

process the connections simultaneously.

Brief Description of the Drawings

Fig. 1 illustrates a simplified graphical representation of a workload
5 sharing mechanism for a multiple processor computer system according to one
example of the present invention.

Fig. 2 illustrates a block diagram for implementing a multiple timer thread
parallel processing system according to another example of the present
invention.

10

Description of the Preferred Embodiment

According to various examples of the present invention, the load
balancing is performed using hash buckets. In essence, the connections are
mapped to hash buckets using a hash function, and then the hash buckets are
15 distributed across multiple processors based on the processor utilization
variances.

Referring now to Fig. 1, a network workload sharing mechanism of a
multiple processor system is shown according to one example of the present
invention. In this configuration, the network 10 is the main channel or the
20 backbone for all information exchange. For illustration purpose only, it is
assumed that the network uses the TCP/IP protocol, IP packets travel
throughout the network. The packets are then fed into a paralellizer 12, which
performs a hash function that maps each packet into a particular hash bucket 15.
A hash function (or a hashing algorithm) is a process of transforming a message
25 of any length into a fixed length output message. The output message is called a
digest or bucket. A hash function can be designed so that it is computationally
unfeasible to determine an input message from its corresponding output bucket.

Although there are significant numbers of hashing algorithms known in the art, for the purpose of the present invention, only those fast or efficient hashing algorithms such as the "CRC based hashing algorithm" are desired. The primary goal for the hashing algorithm used in the examples of the present invention is to
5 distribute a number of packets from different connections across the generated hash buckets. In essence, the hashing algorithm helps to map a non-finite number of connections to a finite number of hash buckets so that they can be controlled or manipulated better.

Not only do packets need to be directed to particular hash buckets, it must
10 also be assured that packets from a predetermined connection go only to a dedicated hash bucket. By this arrangement, packets from a connection will be processed exclusively by a single processor, thereby avoiding unsynchronized processing of information from the same connection. It is well known that the four basic parameters (also known as "4-tuple") of the header section of the
15 packet (i.e., <source address, source port, destination address, destination port>) are unique for each TCP/IP connection and are always identical with all the packets of a particular connection. In one example of the present invention, the hash bucket is identified by using the 4-tuple as a key to perform the hash function. Consequently, it is assured that packets from a particular connection
20 will only be assigned to predetermined bucket.

Further, the same hash bucket can be reused by different connections. That is, packets from different connections can be mapped to the same hash bucket. Moreover, each processor can process multiple hash buckets. As a graphical representation, two or three hash buckets 13 adhere to each other in
25 Fig. 1 to indicate that they are different hash buckets but are destined to the same processor.

Once the parallelizer generates the hash buckets, the packets in the hash

buckets may be queued. Multiple queues are implemented to accommodate the configuration of multiple processors. Each queue has a predetermined number of hash buckets determined by a load balancer 14. The load balancer 14 positions itself between the queues of hash buckets and the processors (e.g., 16a-16n), and
5 is responsible for monitoring the load of each processor and assigning different hash buckets to the processors. The load balancer has the ability to dynamically rearrange the hash buckets in the queues to the processors based on the level of capacity usage of each processor. Therefore, multiple processors can be used intelligently to maximize their collective efficiency.

10 In one example, the processor sequentially takes packets from different hash buckets assigned to it until all the hash buckets are reached in the round. However, it is possible that there are packets from various connections are assigned to a single hash bucket, and it is also likely that the numbers of packets from different connections may vary significantly. A queuing model may be
15 applied to the hash bucket assigned to a predetermined processor to assure that packets from a single connection does not overcrowd (or as it is known, "hog") the processor while other packets from other connections are waiting to be processed by the same processor.

One example of the present invention is further explained below in the
20 context of a four-processor system and a predetermined hash function with 4096 hash buckets. A hash function with 4096 hash buckets simply means that the hash function will generate a fixed-length bucket of 12 bits long. Hence, the values of buckets vary from 0 to 4095, which are used as identification numbers for the corresponding hash buckets. It is also understood that based on the hash
25 function selected, the output can be any number, and doesn't have to be between 0 and 4095.

If it is found that statistically, a processor utilization distribution rate in

terms of the percentage of capacity usage per processor is 10%, 15%, 11%, and 28% respectively for the four processors at a given instant, the total utilization rate for the entire four-processor system is considered to be 64% by summing all individual processor utilization distribution rates. An imbalance distribution index for each processor can be calculated by subtracting the processor utilization distribution rate from the total utilization rate. That is, the four imbalance distribution indexes for the processors are 54%, 49%, 53%, and 36%. A total imbalance index can be calculated by adding all the four imbalance distribution indexes together. In this case, the total imbalance index is (54+49+53+36)=192 (if the percentage sign is deemed to be taken away). The hash buckets are then distributed based on the ratio between the imbalance distribution index of each processor and the total imbalance index. Therefore, the number of hash buckets distributed equitably in this example should be as follows:

Processor 1	$(54/192*4096)=1152,$
Processor 2	$(49/192*4096)=1054,$
Processor 3	$(53/192*4096)=1130,$ and
Processor 4	$(36/192*4096)=763.$

Assuming that a good hash function is incorporated in the multiple processor system, and there are a large number of hash buckets, the load balancing method described above thus proactively distributes the workload. Also since it is based on load imbalance across the processors, the distribution is dynamically adjusted or adapted to the current load situation.

The general concept derived from above-described invention can be used in many areas of a multiprocessor computer system. For example, in a multiple processor system such as Novell's Netware Multiple-Processor system, the TCP/IP workload is distributed across the processors at the connection level.

Although there are multiple timer threads like Delayed Ack Timer and KeepAlive Timer in some implementations, but most of the time, there is only one timer thread running to process a TCP connection table which lists all active TCP connections. The existing solutions in all implementation for TCP timers do
5 not utilize the parallel execution of timer threads on different processors. In conventional methods, a single timer thread sequentially processes each connection one after another. The average delay between timer processing is likely to increase in proportion to the number of connections.

The present embodiment provides a unique way of splitting the process of
10 handling the TCP Connection Table by the timer across multiple processors by mapping the connections in the TCP Connection Table to multiple hash buckets. According to one example of the present invention, at any instant of time a group of connections from the TCP Connection Table go through a hash function and are further mapped to each participating processor according to the result of
15 the hash function. This mechanism allows the use of multiple timer threads simultaneously, and each timer thread running on a specific processor and processing disjoint sets of TCP connections at any instant of time.

Referring now to Fig. 2, a block diagram is shown for describing a use of multiple timer threads in a multiple processor system. Instead of using just a
20 single timer thread, it is now possible to use multiple timer threads and each timer thread can process multiple hash buckets individually (e.g., N number timer threads with each processing on $4096/N$ buckets, where N is the number of processors). Since there is a single timer thread for each processor, and it processes a group of connections assigned to that particular processor, no
25 connection can be processed by more than one timer thread. Hence, there is no need to have any locking mechanism for accessing connections in the Connection Table. As it is well known, a locking mechanism is needed only when more than

one timer thread is accessing the Connection Table simultaneously. In terms of scalability, in case there are more than one processor and a large number of connections, employing multiple timer threads can be a true scalable solution. And in the immediate example, it can be said that the method employing N
5 number timer threads is N times scalable than the conventional method.

As shown in Fig. 2, it is assumed that there are altogether 4096 hash buckets in a four-processor system and with timer threads scheduled at a fixed time interval. In this configuration, each TCP connection is mapped to one of these buckets based on four header parameters (i.e., local address, local port,
10 remote address, and remote port) which are unique for each TCP/IP connection. Each hash bucket is further mapped to one of the participating processors (e.g., processor 0 to 3) based on the current workload of each processor. In this configuration, it is predetermined that four timer threads are spawned on the four processors every 224 milliseconds to process active TCP connections. For
15 instance, the timer thread running on processor 0 first processes a set of connections 20 (which contains connections C1, C3, C15, and C16) mapped to the first hash bucket 22 (which is identified by its value "0"). Thereafter, the timer thread for processor 0 continues to process another set of connections 24 as identified by the hash bucket 2 (numeral 26). Apart from the timer thread for
20 processor 0, other timer threads are working for other processors.

As shown in Fig. 2, the hash bucket 1 (numeral 28) is assigned to processor 1, and For example, it brings another set of connections 30 (i.e., C2, C5, and C11) for processor 1 to operate on. In the same fashion, the other two timer threads can independently and simultaneously work for processors 3 and 4 for
25 processing connections. It can be comfortably assumed that, in a multiple processor system, a good hash function can be selected and a large number of buckets for mapping the connections are available as described earlier.

One advantage of the present invention is that it makes it possible to process the Connection Table by multiple timer threads in a parallel fashion and it also eliminates the need for a locking mechanism. The example of the present invention thus assures a faster processing of a collection of connections such as a
5 Connection Table on a multiple processor system without resource contention between timer threads. Further, the present invention is scalable with regard to the number of processors and number of connections.

The above disclosure provides many different embodiments, or examples, for implementing different features of the invention. Specific examples of
10 components, and processes are described to help clarify the invention. These are, of course, merely examples and are not intended to limit the invention from that described in the claims. For example, various acceptable hashing algorithms can also be used by one skilled in the art.

While the invention has been particularly shown and described with
15 reference to the preferred embodiment thereof, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention, as set forth in the following claims.

WHAT IS CLAIMED IS:

1 1. A method for balancing a workload for a plurality of processors in
2 a multiple processor computer system, the system designed for processing a
3 plurality of packets from a plurality of connections, the method comprising:
4 assigning a packet to a hash bucket determined by performing a
5 predetermined hash function; and
6 queuing the hash bucket to a processor so that the workload of all the
7 processors are balanced,
8 wherein the hash function relates to a predetermined number of hash
9 buckets, and wherein a plurality of packets from different connections can be
10 assigned to the same hash bucket.

1 2. The method of claim 1 wherein the step of assigning includes a step
2 of using a source address, source port, destination address, and destination port
3 parameters from the header section of the packet to identify the hash bucket.

1 3. The method of claim 1 wherein the hash function is designed to
2 assign a plurality of packets from a connection to the same hash bucket.

1 4. The method of claim 1 further comprising monitoring the workload
2 of each processor involved in the system.

1 5. The method of claim 1 wherein the step of queuing further
2 includes:
3 identifying statistically a percentage of capacity usage for each processor;
4 obtaining imbalance distribution index for each processor from the
5 identified percentage of capacity usage;
6 distributing the hash buckets in proportion to the imbalance distribution
7 index and in reverse proportion to a total imbalance index,
8 wherein the total imbalance index is a summation of all obtained
9 imbalance distribution indexes.

1 6. The method of claim 1 wherein the step of queuing further includes
2 a step of applying a queuing model for packets in the hash bucket to assure that
3 packets from a particular connection do not unduly occupy the processor for an
4 undesirable time period.

1 7. A computer program for balancing workload for a plurality of
2 processors in a multiple processor computer system, the system designed for
3 processing a plurality of packets from a plurality of connections, the computer
4 program comprising:
5 means for assigning a packet to a hash bucket determined by performing a
6 predetermined hash function; and
7 means for queuing the hash bucket to a processor so that the workload of
8 all the processors are balanced,
9 wherein the hash function has a predetermined number of hash buckets,
10 and wherein a plurality of packets from different connections can be assigned to
11 the same hash bucket.

1 8. The computer program of claim 7 wherein the means for assigning
2 includes means for using a source address, source port, destination address, and
3 destination port parameters from the header section of the packet to identify the
4 hash bucket.

1 9. The computer program of claim 7 wherein the hash function is
2 designed to assign a plurality of packets from a connection to the same hash
3 bucket.

1 10. The computer program of claim 7 further comprising means for
2 monitoring the workload of each processor involved in the system.

1 11. The computer program of claim 7 wherein the means for queuing
2 further includes means for:
3 identifying statistically a percentage of capacity usage for each processor;
4 obtaining imbalance distribution index for each processor from the
5 identified percentage of capacity usage;
6 distributing the hash buckets in proportion to the imbalance distribution
7 index and in reverse proportion to a total imbalance index,
8 wherein the total imbalance index is a summation of all obtained
9 imbalance distribution indexes.

1 12. The computer program of claim 7 wherein the means for queuing
2 further includes means for applying a queuing model for packets in the hash
3 bucket to assure that packets from a particular connection do not unduly occupy
4 the processor for an undesirable time period.

1 13. A method for processing a plurality of connections with a plurality of
2 timer threads by a plurality of computer processors in a multiple processor
3 computer system, the method comprising:
4 providing a plurality of hash buckets related to a hash function;
5 mapping a connection to one of the hash buckets; and
6 assigning each hash bucket to a processor timer thread based on a
7 workload thereof so that the processor only processes the connection mapped to
8 the assigned hash bucket,
9 wherein a plurality of timer threads for the computer processors thus
10 process a plurality of connections simultaneously.

1 14. The method of claim 13 further comprising monitoring the
2 workload of the each processor.

1 15. The method of claim 13 wherein the step of mapping is based on four
2 connection parameters, which are local address, local port, remote address, and
3 remote port.

1 16. A computer program for processing a plurality of connections with a
2 plurality of timer threads by a plurality of computer processors in a multiple
3 processor computer system, the program comprising:

4 a plurality of hash buckets related to a predetermined hash function;

5 means for mapping a connection to one of the hash buckets; and

6 means for assigning each hash bucket to a processor timer thread based on
7 a workload thereof so that the processor only processes the connection mapped
8 to the assigned hash bucket,

9 wherein a plurality of timer threads for the computer processors thus
10 process a plurality of connections simultaneously.

1 17. The program of claim 16 further comprising means for monitoring
2 the workload of the each processor.

1 18. The program of claim 16 wherein the means for mapping is based
2 on four connection parameters, which are local address, local port, remote
3 address, and remote port.

SYSTEM AND METHOD FOR BALANCING WORKLOAD BASED ON HASH BUCKETS

Abstract

A method and system is provided for balancing work load for a plurality of computer processors in a multiple processor computer system. In such a system, there are a plurality of packets to be processed. First, a plurality of hash
5 buckets are generated and each packet is assigned to a bucket by performing a predetermined hash function. The different hash buckets are then assigned to the processors based on the workload of all the processors. In connection with processing connections with a timer thread, a plurality of connections are processed with a plurality of timer threads. A plurality of hash buckets are
10 provided, and each connection is mapped to one of the hash buckets. Further, each hash bucket is assigned to a processor timer thread based on a workload thereof.

**SYSTEM AND METHOD FOR BALANCING TCP/IP/WORKLOAD OF
MULTI-PROCESSOR SYSTEM BASED ON HASH BUCKETS**

Inventors: Vipul Shah
1048 S. Winchester Blvd
Apt #12
San Jose, CA 95128
Citizenship: India

N. S. S. Narasimha Rao
1-370, Rayudupalem
Kakinada, India
Citizenship: India

Alka Agrawal
310, Mukund Apartment
Victoria Layout
Bangalore, Karnataka
Citizenship: India

Subrata Sarkar
M6, Royal Castle, Kodihalli
HAL IIIrd Stage
Bangalore 560008
India
Citizenship: India

Kumar Subramanian
1521, 2nd Floor
2nd Cross, 13th Main
HAL 2nd Stage, Kodihalli
Bangalore - 560008
Citizenship: India

Himanshu Shukla
380, 1st Floor
1st A Cross,
7th Block, Kormangala
Bangalore - 560095
Citizenship: India

Assignee: Novell, Inc.
122 East 1700 South
Provo, Utah 84606

HAYNES AND BOONE, LLP
901 Main Street, Suite 3100
Dallas, Texas 75202-3789
(214) 651-5000
Attorney Docket No. 26530.23




Fig. 1

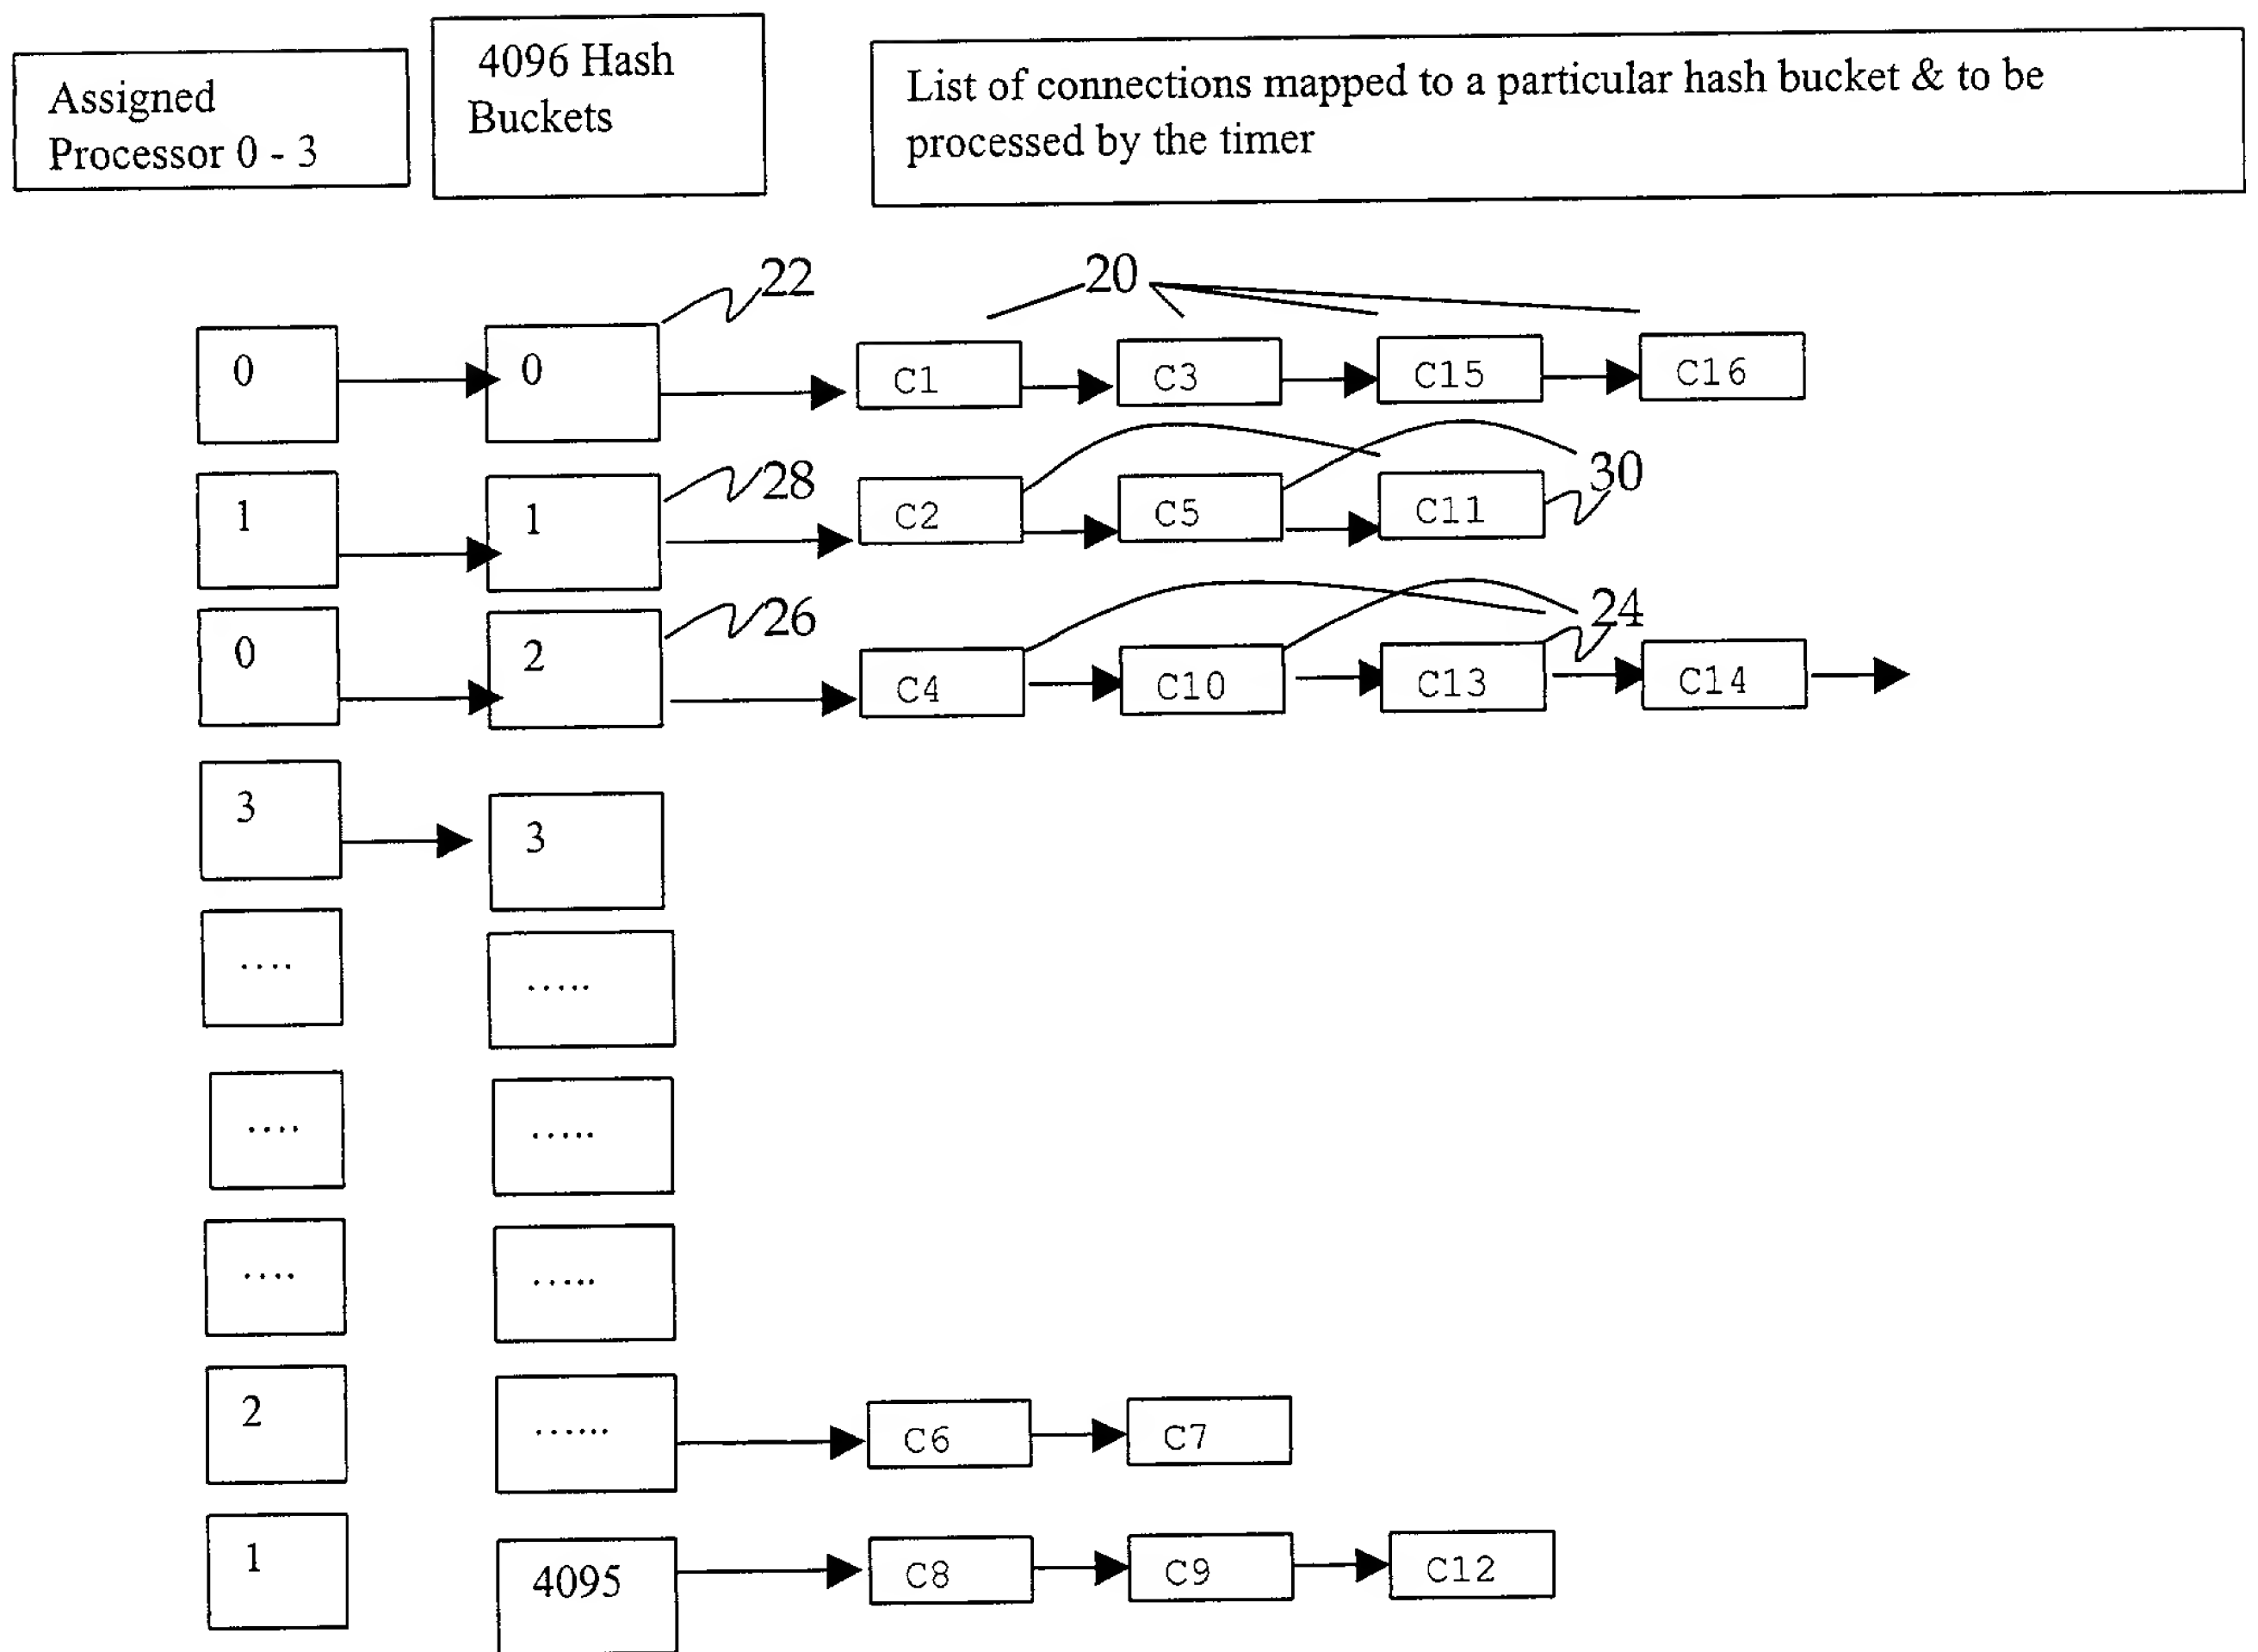


Fig. 2

**DECLARATION AND POWER OF ATTORNEY FOR
PATENT APPLICATION**

As below named inventor, I hereby declare that:

My residence, post office address and citizenship is as stated below next to my name;

I believe I am the original and first of the subject matter which is claimed and for which a patent is sought on the invention entitled

**SYSTEM AND METHOD FOR BALANCING TCP/IP WORKLOAD OF
MULTI-PROCESSOR SYSTEM BASED ON HASH BUCKETS**

the specification of which: (check one)

XXX is attached hereto.

_____ was filed on _____
under Attorney's Docket Number 26530.23 (IDR-464/5)
as Application Serial No. _____
and was amended on _____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 CFR 1.56.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 USC 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby claim the benefit under Title 35, United States Code, § 119(e) of any United States provisional application(s) listed below.

(Application Number)

(Filing Date)

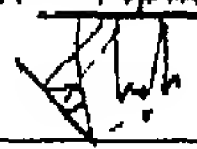
DOCKET NO.: 26530.23 (IDR-464/5)

POWER OF ATTORNEY: As named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Theodore Baroody	Reg. No. 45,417	Warren B. Kice	Reg. No. 22,732
Jeffrey M. Becker	Reg. No. 35,442	Christopher R. Kosh	Reg. No. 42,760
James R. Bell	Reg. No. 26,528	Michael J. Balconi-Lamica	Reg. No. 34,291
Daniel E. Burke	Reg. No. P-46,588	Todd Mattingly	Reg. No. 40,298
Michael S. Bush	Reg. No. 31,745	David L. McCombs	Reg. No. 32,271
L. Howard Chen	Reg. No. P-46,615	Bill R. Naifeh	Reg. No. 44,962
Randall E. Colson	Reg. No. 40,566	David M. O'Dell	Reg. No. 42,044
Michael A. Davis, Jr.	Reg. No. 35,488	Phillip B. Philbin	Reg. No. 35,979
Ruben C. DeLeon	Reg. No. 37,812	Constance M. Pielech	Reg. No. P46,991
Timothy Headley	Reg. No. 31,765	Brandi W. Sarfatis	Reg. No. 37,713
Brian J. Hubbard	Reg. No. 45,873	David O. Simmons	Reg. No. 43,124
Rita M. Irani	Reg. No. 31,028		

Send correspondence to David L. McCombs, Haynes and Boone, LLP, 901 Main Street, Suite 3100 Dallas, Texas 75202-3789 and direct all telephone calls to David L. McCombs at 214/651-5533.

Full Name of First Inventor: Vipul Shah

Inventor's Signature: 

Dated: 10/4/00

Residence: 1048 S. Winchester Blvd, Apt. 12, San Jose, California 95128

Citizenship: India

Post Office Address: 1048 S. Winchester Blvd, Apt. 12, San Jose, California 95128

Full Name of Second Inventor: N.S.S. Narasimha Rao

Inventor's Signature: _____

Dated: _____

Residence: 1-370, Rayudupalem, Kakinada, India

Citizenship: India

Post Office Address: 1-370, Rayudupalem, Kakinada, India

Full Name of Third Inventor: Alka Agrawal

POWER OF ATTORNEY: As named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Theodore Baroody	Reg. No. 45,417	Warren B. Kice	Reg. No. 22,732
Jeffrey M. Becker	Reg. No. 35,442	Christopher R. Kosh	Reg. No. 42,760
James R. Bell	Reg. No. 26,528	Michael J. Balconi-Lamica	Reg. No. 34,291
Daniel E. Burke	Reg. No. P-46,588	Todd Mattingly	Reg. No. 40,298
Michael S. Bush	Reg. No. 31,745	David L. McCombs	Reg. No. 32,271
L. Howard Chen	Reg. No. P-46,615	Bill R. Naifeh	Reg. No. 44,962
Randall E. Colson	Reg. No. 40,566	David M. O'Dell	Reg. No. 42,044
Michael A. Davis, Jr.	Reg. No. 35,488	Phillip B. Philbin	Reg. No. 35,979
Ruben C. DeLeon	Reg. No. 37,812	Constance M. Pielech	Reg. No. P46,991
Timothy Headley	Reg. No. 31,765	Brandi W. Sarfatis	Reg. No. 37,713
Brian J. Hubbard	Reg. No. 45,873	David O. Simmons	Reg. No. 43,124
Rita M. Irani	Reg. No. 31,028		

Send correspondence to David L. McCombs, Haynes and Boone, LLP, 901 Main Street, Suite 3100 Dallas, Texas 75202-3789 and direct all telephone calls to David L. McCombs at 214/651-5533.

Full Name of First Inventor: Vipul Shah

Inventor's Signature: _____ Dated: _____

Residence: 1048 S. Winchester Blvd, Apt. 12, San Jose, California 95128

Citizenship: India

Post Office Address: 1048 S. Winchester Blvd, Apt. 12, San Jose, California 95128

Full Name of Second Inventor: N.S.S. Narasimha Rao

Inventor's Signature: N.S.S. Narasimha Rao Dated: 08/10/2000

Residence: 1-370, Rayudupalem, Kakinada, India

Citizenship: India

Post Office Address: 1-370, Rayudupalem, Kakinada, India

Full Name of Third Inventor: Alka Agrawal

DOCKET NO.: 26530.23 (IDR-464/5)

Full Name of Third Inventor: Alka Agrawal

Inventor's Signature:  Dated: 8th Oct 2000

Residence: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Citizenship: India

Post Office Address: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Full Name of Fourth Inventor: Subrata Sarkar

Inventor's Signature: Subrata Sarkar Dated: 5/10/5/2000

Residence: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Citizenship: India

Post Office Address: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Full Name of Fifth Inventor: Kumar Subramanian

Inventor's Signature: _____ Dated: _____

Residence: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli, Bangalore 56008

Citizenship: India

Post Office Address: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli,

Bangalore 56008

DOCKET NO.: 26530.23 (IDR-464/5)

Full Name of Third Inventor: Alka Agrawal

Inventor's Signature: _____ Dated: _____

Residence: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Citizenship: India

Post Office Address: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Full Name of Fourth Inventor: Subrata Sarkar

Inventor's Signature: _____ Dated: _____

Residence: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Citizenship: India

Post Office Address: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Full Name of Fifth Inventor: Kumar Subramanian

Inventor's Signature:  Dated: 9th October 2000

Residence: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli, Bangalore 56008

Citizenship: India

Post Office Address: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli,

Bangalore 56008

DOCKET NO.: 26530.23 (IDR-464/5)

Inventor's Signature: _____ Dated: _____

Residence: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Citizenship: India

Post Office Address: 310, Mukund Apartment, Victoria Layout, Bangalore, Karnataka

Full Name of Fourth Inventor: Subrata Sarkar

Inventor's Signature: _____ Dated: _____

Residence: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Citizenship: India

Post Office Address: M6, Royal Castle, Kodihalli, HAL IIIrd Stage, Bangalore 560008, India

Full Name of Fifth Inventor: Kumar Subramanian

Inventor's Signature: _____ Dated: _____

Residence: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli, Bangalore 56008

Citizenship: India

Post Office Address: 1521 2nd Floor, 2nd Cross, 13th Main, HAL 2nd Stage, Kodihalli,

Bangalore 56008

Full Name of Sixth Inventor: Himanshu Shukla

Inventor's Signature:  Dated: 06/10/2000

Residence: 380, 1st Floor, 1st A Cross, 7th Block, Kormangala, Bangalore 560095

Citizenship: India

Post Office Address: 380, 1st Floor, 1st A Cross, 7th Block, Kormangala, Bangalore 560095

D-818090.1